



Improving Backup Performance with Defragmentation (A business application study) June, 2006

Executive Summary:

With continued expansion of organizational storage allocation requirements, the time required for file-based data backup also continues to increase yet the time allotted for this data backup remains constant, exposing a critical issue. Though backup software vendors continue making advancements in technology to speed up the process, disk fragmentation will continue to remain an unresolved issue that must be addressed on the file system level to solve this bottleneck.

In business environments, backing up local and network data is commonly done with both tape devices and disk devices. As data backup involves file access, fragmentation of data files can have a profound impact on the length of time a backup procedure may take.

Diskeeper Corporation, an industry-leading developer of systems management tools, performed a series of rigorous tests to determine the effects of fragmentation on backup operations and the benefits of defragmentation.

In conclusion, Diskeeper Corporation has found that the time required to back up data from a typical hard drive volume to these backup devices can be *decreased* and the backup data transfer rate *increased* by defragmenting the volume prior to backup. Furthermore, directory consolidation (defragmenting and consolidating directory files) has a direct additional benefit in both of these areas. The study showed improvements in backup times up to 69% (decrease in total back up time)—*less than one half* of the time for the fragmented case. Similarly, backup data transfer rates increased up to 69% over that for the fragmented state—nearly 1.7 as fast.

Test Methodology:

Overview

A set of trials was conducted to determine the benefit that file and directory defragmentation have on file backup performance for a server.

Setup

Hardware:

1. Local Test system running Windows 2003 Server Enterprise with Service Pack 1:
 - Intel D925CV2 motherboard (BIOS date 10/12/04) with 3.6 GHz Pentium 4 CPU
 - 2GB RAM
 - 200GB Seagate ST3200826AS SATA hard drive

2. Network remote backup system running Windows 2003 Server Enterprise with Service Pack 1:
 - Dual 1GHz Pentium 3 CPUs on an Asus CUR-DLS motherboard, BIOS revision 1009
 - 1GB ECC memory (4x256MB DIMMs)
 - Seagate 36GB 10,000 RPM Cheetah SCSI hard drive
 - Mylex Extreme RAID 2000 SCSI-160 controller with 32MB cache memory, BIOS PLUS version 6.01-39.
 - Intel PRO/1000 MT 1Gbit/s network interface card
 - Adaptec 2940-UW SCSI controller (for local tape backup system)
 - Exabyte EZ-17 Mammoth-2 tape external backup system (connected to outboard 69-pin connector on Adaptec 2940 via 6 foot shielded cable)

3. Network remote backup system running Windows 2000 Server with Service Pack 4:
 - Dual 1GHz Pentium 3 CPUs on a DL380 motherboard, BIOS revision 2.10 (revised 12/18/02)
 - 1GB DDR (2 x 512MB DIMMs)
 - Seagate ST380011A IDE 80GB hard drive

Backup Software Used:

1. Symantec Backup Exec 10d (Build 5629)
2. Computer Associates BrightStor ARCserve Backup release 11.5 (Build 3884)

Defragmentation Software Used:

Diskeeper® 10.0.

Test methodology:

1. Create an image of the operating system partition on the test server.
2. Create and image a fragmented 20 gigabyte test volume containing approximately 100,000 files, about 60% data fragmentation and about 40% free space. (See *Figure 1* in the Appendix.)
3. Install and set up one of the backup programs to test with (Backup Exec or ARCserve Backup).
4. Back up the test volume to a local disk backup device (created with the backup program), and record the backup time and data transfer rate for the fragmented volume.
5. Install the backup software on another system on the same network.
6. Back up the test volume over the network to a disk backup device (created with the backup software), and record the backup time and data transfer rate for the fragmented volume.
7. Back up the test volume over the network to a tape backup device, and record the backup time and data transfer rate over the fragmented volume.¹
8. Defragment the test volume using Diskeeper Server Enterprise Edition (using different defragmentation methods to see what the results were).
9. Repeat steps 4-8 with the defragmented test volume.
10. Restore the image of the fragmented test volume.
11. Repeat steps 3-9 for the next backup program.

¹ The remote testing of backing up the test volume to a network tape backup device was done only with the Backup Exec program.

Analysis:

Table 1 summarizes the results obtained in backing up the 20GB test volume to a local disk device, a network tape backup device and a network disk device using Backup Exec software, and a local disk device and a network disk device using ARCserve Backup software. The results before and after defragmentation with Diskeeper 10.0 Server Enterprise Edition are shown.

In all tests run it was observed that defragmenting the test volume prior to backing it up reduced the backup time and increased the backup data transfer rate. The results of many tests indicate directory consolidation has a large impact on backup processes when backing up to network disk backup devices.

The data transfer rates for Backup Exec trials improved from 6% to 16%, between Network disk backup trials and local disk backup trials.

ARCserve Backup's trials benefited similarly from defragmentation, the backup data transfer rates increasing between 2% – 12%, between network disk backup trials and local disk backup trials.

Diskeeper's defragmentation mode was set to Max Disk Performance and a manual defragmentation run was done for the trials for both Backup Exec and ARCserve Backup. This produced the results shown in Figure 2 in the Appendix. Because the directory files were scattered across the volume, a set of trials was performed with both backup software products to see if a consolidation of these directories with a boot-time defragmentation would improve performance. These results can be seen in the third column in Table 1 above (see also Figure 3 in the Appendix).

The results show that in addition to defragmenting the case, consolidating directories had a further benefit on backup performance. Backup times for ARCserve were improved by an additional 3% for the local disk backup, and an additional 41% for the network disk backup, a marked improvement. The network tape backup done for Backup Exec improved by 6% over the fragmented case after a defragmentation run with directory consolidation. Finally, the network disk backup done with Backup Exec showed a 20% improvement in both the elapsed time and transfer rate after directory consolidation, above and beyond the 6% improvement seen after simple defragmentation.

Table 1. Summary of Backup results before and after defragmentation (from backup log files – see Appendices)

ARCServe Backup

	Fragmented Volume	Defragmented Volume	Defragmented Volume (with Dir Consolidation)	Defragmented Volume (with I-FAAST & Dir Consolidation)
Backup from Test Volume to Local Disk Backup Device				
Time to Backup Test Volume, hr:min:sec	0:05:37	0:04:58	0:04:46	0:04:50
Percent Decreased from fragmented case	→	11.57%	15.13%	13.95%
Data Transfer Rate, MB/min	2,168.74	2,452.57	2,555.48	2,521.03
Percent Improvement over fragmented case	→	11.57%	15.13%	13.97%
Backup from Test Volume to Network Disk Backup Device				
Time to Backup Test Volume, hr:min:sec	0:16:46	0:16:18	0:09:24	0:05:07
Percent Decreased from fragmented case	→	2.78%	43.94%	69.48%
Data Transfer Rate, MB/min	726.5	747.3	1,295.86	2,381.43
Percent Improvement over fragmented case	→	11.57%	43.94%	69.49%

Backup Exec

	Fragmented Volume	Defragmented Volume	Defragmented Volume (with Dir Consolidation)	Defragmented Volume (with I-FAAST & Dir Consolidation)
Backup from Test Volume to Local Disk Backup Device				
Time to Backup Test Volume, hr:min:sec	0:06:16	0:05:20	0:05:37	0:05:23
Percent Decreased from fragmented case	→	14.89%	10.37%	14.10%
Data Transfer Rate, MB/min	1,991.00	2,368.00	2,235.00	2,322.00
Percent Improvement over fragmented case	→	15.92%	10.92%	14.25%
Backup from Test Volume to Network Disk Backup Device				
Time to Backup Test Volume, hr:min:sec	0:08:50	0:08:17	0:06:23	0:06:11
Percent Decreased from fragmented case	→	6.23%	27.74%	30.00%
Data Transfer Rate, MB/min	1,386.00	1,478.00	1,942.00	1,991.00
Percent Improvement over fragmented case	→	6.22%	28.63%	30.39%
Backup from Test Volume to Network Tape Backup Device				
Time to Backup Test Volume, hr:min:sec	0:09:44	0:09:15	0:08:30	0:08:31
Percent Decreased from fragmented case	→	4.97%	12.67%	12.50%
Data Transfer Rate, MB/min	1,221.00	1,284.00	1,398.00	1,395.00
Percent Improvement over fragmented case	→	4.91%	12.66%	12.47%

Analysis continued:

Combining the benefits of directory consolidation from boot-time defragmentation with the I-FAAST™ (Intelligent File Access Acceleration Sequencing Technology) engine improved backup data transfer rates and backup times in the network disk backup trials for ARCserve by a total of 70%—far less than half the time to back up the fragmented case. Backup Exec showed 30% improvement over the fragmented case with defragmentation and boot-time directory consolidation with I-FAAST.

Network tape backup trials performed using Backup Exec showed gradient improvement upon the fragmented case, defragmenting the volume prior to backup improved backup time and data transfer rate by 5%, while going on to perform a directory consolidation (also with I-FAAST) improved the backup time and data transfer rate by nearly 13% over the original fragmented case.

Diskeeper Corporation

Appendix:

1. 20GB test case properties before defragmentation:
(summary from Diskeeper Server Enterprise Edition 10.0.597)

Volume Files

Volume size	= 20,480 MB
Cluster size	= 4 KB
Used space	= 12,270 MB
Free space	= 8,209 MB
Percent free space	= 40 %

Fragmentation percentage

Volume fragmentation	= 37 %
Data fragmentation	= 61 %

Directory fragmentation

Total directories	= 21
Fragmented directories	= 8
Excess directory fragments	= 3,175

File fragmentation

Total files	= 100,000
Average file size	= 122 KB
Total fragmented files	= 20,887
Total excess fragments	= 1,859,502
Average fragments per file	= 19.59
Files with performance loss	= 20,887

Paging file fragmentation

Paging/Swap file size	= 0 bytes
Total fragments	= 0

Master File Table (MFT) fragmentation

Total MFT size	= 117 MB
MFT records In Use	= 100,056
Percent MFT in use	= 83 %
Total MFT fragments	= 3

Appendix continued:

- Graphical representation of the Test Case before Defragmentation (from Diskeeper 10.0.597):

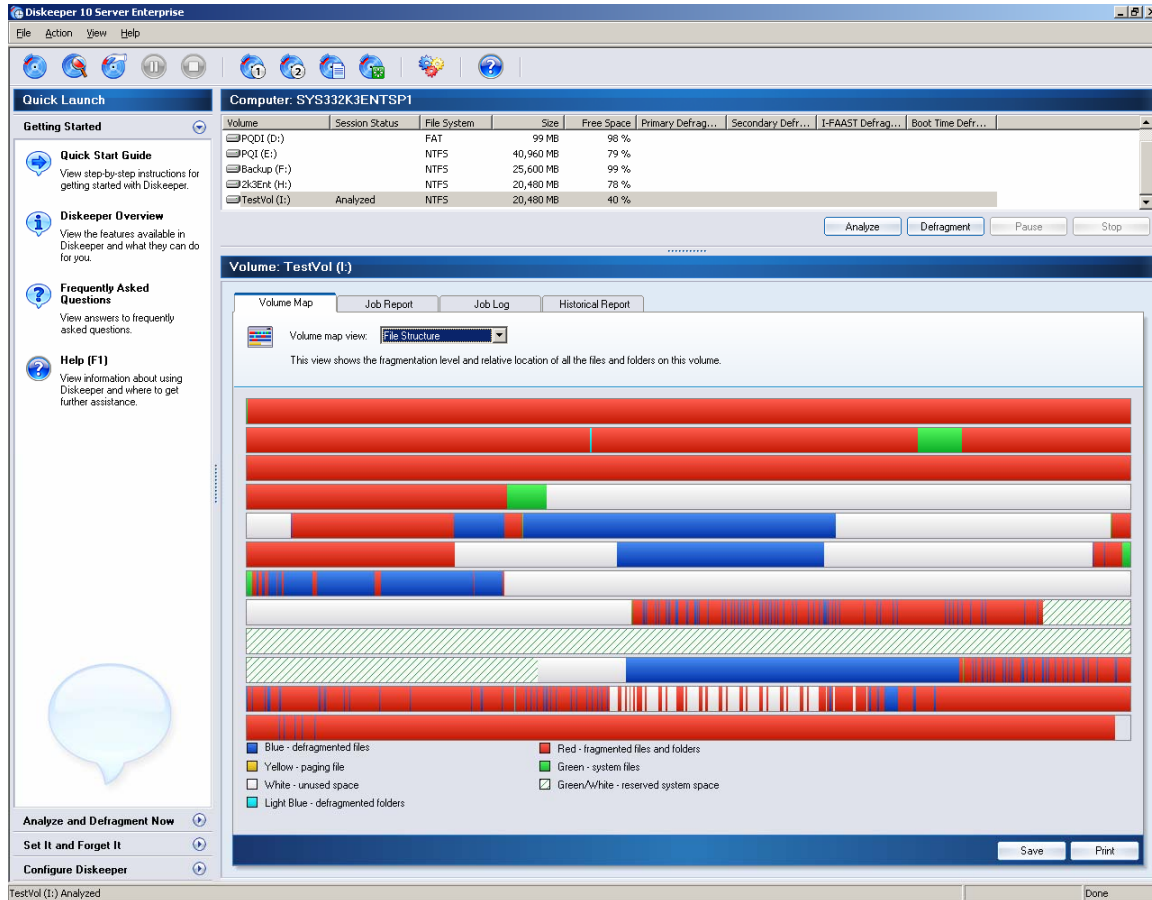


Figure 1. 20GB Test Volume prior to defragmentation.

Appendix continued:

- Graphical representation of the Test Case after manual Defragmentation with default settings (from Diskeeper 10.0.597):

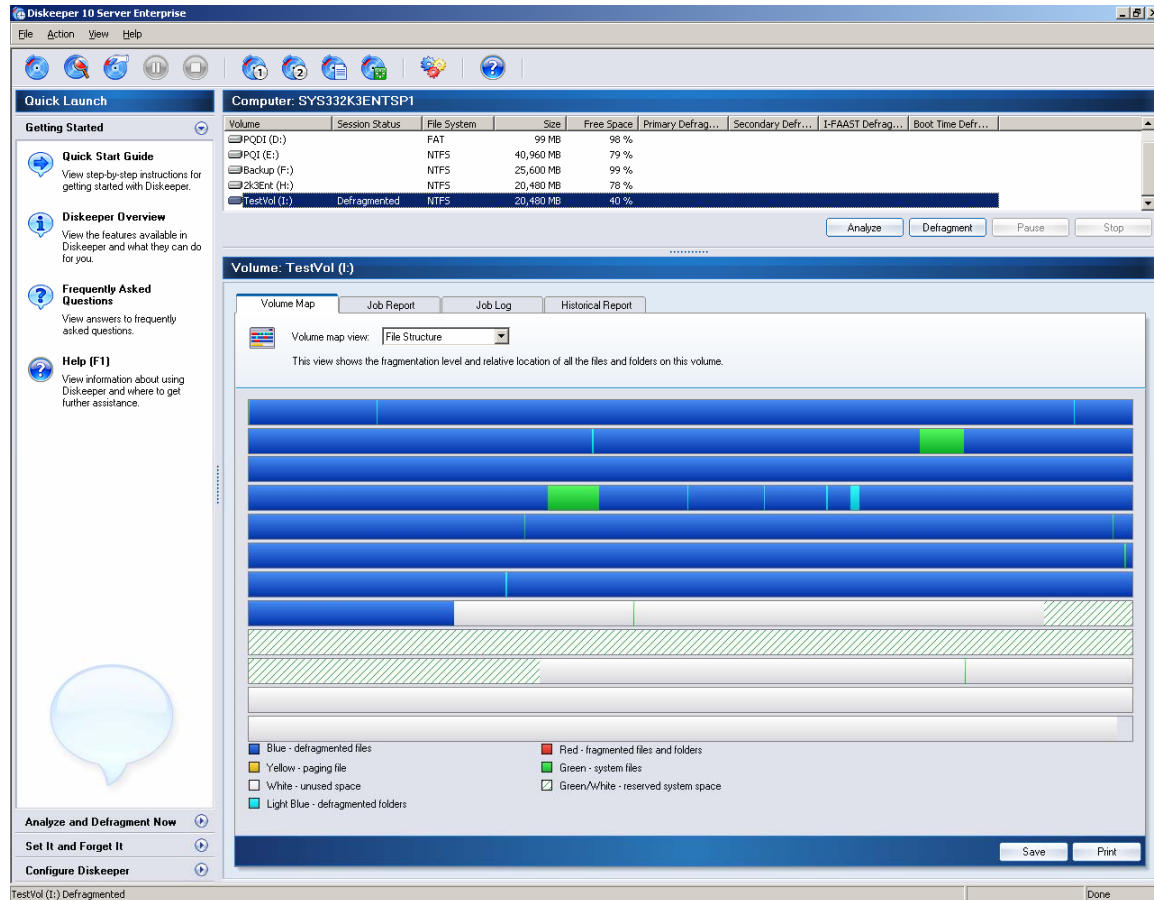


Figure 2. 20GB Test Volume after one run manual defragmentation set for Max Disk Performance (default setting).

Appendix continued:

4. Graphical representation of the Test Case after Boot-time Defragmentation (from Diskeeper 10.0.597):

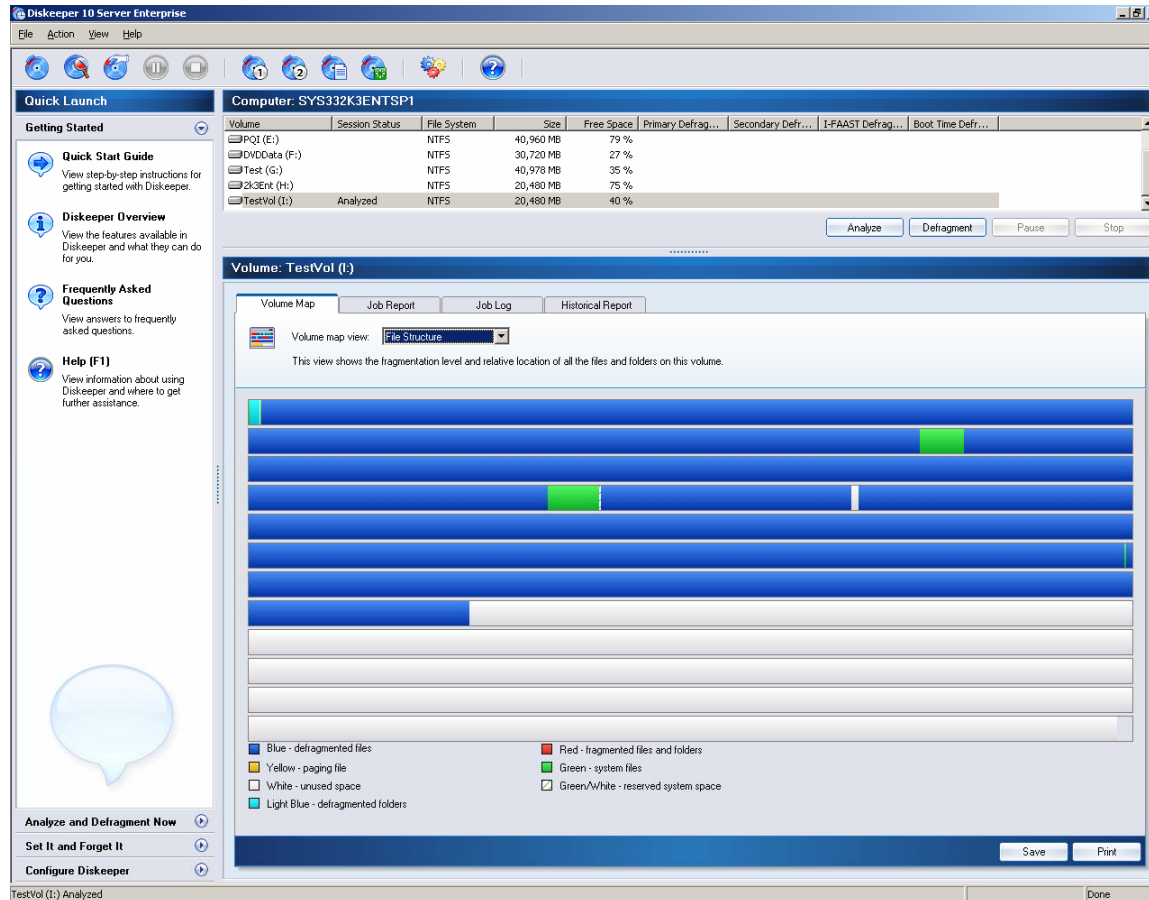


Figure 3. 20GB Test Volume after Boot-time defragmentation to consolidate Directories and Defragment the MFT.

Appendix continued:

5. Backup Exec 10d: Test volume to local disk backup device before defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : 1:41:12 PM
End time             : 1:47:28 PM
Files                : 99,996
Directories          : 20
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,457,438,885 bytes
Rate                 : 1,991.00 MB/Min
```

6. Backup Exec 10d: Test volume to local disk backup device after defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : 2:12:18 PM
End time             : 2:17:38 PM
Files                : 99,996
Directories          : 20
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,457,438,885 bytes
Rate                 : 2,368.00 MB/Min
```

7. Backup Exec 10d: Test volume to local disk backup device after boot-time defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : Monday, May 01, 2006 1:23:06 PM
End time             : Monday, May 01, 2006 1:28:43 PM
Files                : 99,997
Directories          : 21
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,461,502,481 bytes
Rate                 : 2,235.00 MB/Min
```

Appendix continued:

8. Backup Exec 10d: Test volume to local disk backup device after IFAAST engine run results (excerpted from the backup log):

```
Set status           : Completed
Resource name       : \\SYS332K3ENTSP1\I:
Agent used          : Yes
Start time          : Thursday, April 27, 2006 3:23:45 PM
End time            : Thursday, April 27, 2006 3:29:20 PM
Files               : 100,000
Directories         : 22
Skipped files       : 0
Corrupt files       : 0
Files in use        : 0
Byte count          : 12,461,504,506 bytes
Rate                : 2,285.00 MB/Min
```

9. Backup Exec 10d: Test Volume to local disk backup device after boot-time defragmentation & IFAAST engine run results (excerpted from the backup log):

```
Set status           : Completed
Resource name       : \\SYS332K3ENTSP1\I:
Agent used          : Yes
Start time          : Tuesday, May 02, 2006 10:20:37 AM
End time            : Tuesday, May 02, 2006 10:26:00 AM
Files               : 99,998
Directories         : 21
Skipped files       : 0
Corrupt files       : 0
Files in use        : 0
Byte count          : 12,461,503,885 bytes
Rate                : 2,322.00 MB/Min
```

10. Backup Exec 10d: Test volume to network disk backup device before defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name       : \\SYS332K3ENTSP1\I:
Agent used          : Yes
Start time          : 1:14:34 PM
End time            : 1:23:24 PM
Files               : 99,996
Directories         : 20
Skipped files       : 0
Corrupt files       : 0
Files in use        : 0
Byte count          : 12,457,438,885 bytes
Rate                : 1,386.00 MB/Min
```

Appendix continued:

11. Backup Exec 10d: Test volume to network disk backup device after defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : Tuesday, April 18, 2006 2:23:54 PM
End time             : Tuesday, April 18, 2006 2:32:11 PM
Files                : 99,996
Directories          : 20
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,457,438,885 bytes
Rate                 : 1,478.00 MB/Min
```

12. Backup Exec 10d: Test volume to network disk backup device after boot-time defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : Monday, May 01, 2006 1:56:19 PM
End time             : Monday, May 01, 2006 2:02:42 PM
Files                : 99,997
Directories          : 21
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,461,502,481 bytes
Rate                 : 1,942.00 MB/Min
```

13. Backup Exec 10d: Test volume to network disk backup device after IFAAST engine run results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : Friday, April 28, 2006 9:53:03 AM
End time             : Friday, April 28, 2006 10:01:21 AM
Files                : 100,000
Directories          : 22
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,461,504,506 bytes
Rate                 : 1,467.00 MB/Min
```

Appendix continued:

14.Backup Exec 10d: Test Volume to network disk backup device after boot-time defragmentation & IFAAST engine run results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : Tuesday, May 02, 2006 10:33:01 AM
End time             : Tuesday, May 02, 2006 10:39:12 AM
Files                : 99,998
Directories          : 21
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,461,503,885 bytes
Rate                 : 1,991.00 MB/Min
```

15.Backup Exec 10d: Test Volume to network tape backup device before defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : Thursday, May 25, 2006 10:13:13 AM
End time             : Thursday, May 25, 2006 10:22:57 AM
Files                : 99,998
Directories          : 21
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,461,503,885 bytes
Rate                 : 1,221.00 MB/Min
```

16.Backup Exec 10d: Test Volume to network tape backup device after boot-time defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\SYS332K3ENTSP1\I:
Agent used           : Yes
Start time           : Thursday, May 25, 2006 10:49:47 AM
End time             : Thursday, May 25, 2006 10:58:17 AM
Files                : 99,998
Directories          : 21
Skipped files        : 0
Corrupt files        : 0
Files in use         : 0
Byte count           : 12,461,503,885 bytes
Rate                 : 1,398.00 MB/Min
```

Appendix continued:

17. BrightStor ARCserve Backup 11.5: Test volume to local disk backup device before defragmentation results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 5m 37s
Average Throughput 2,168.74 MB/min
Session Status Finished

18. BrightStor ARCserve Backup 11.5: Test volume to local disk backup device after defragmentation results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 4m 58s
Average Throughput 2,452.57 MB/min
Session Status Finished

19. BrightStor ARCserve Backup 11.5: Test volume to local disk backup device after boot-time defragmentation results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 5m 39s
Average Throughput 2,155.95 MB/min
Session Status Finished

20. BrightStor ARCserve Backup 11.5: Test volume to local disk backup device after IFAAST engine run results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 5m 1s
Average Throughput 2,428.90 MB/min
Session Status Finished

Appendix continued:

21. BrightStor ARCserve Backup 11.5: Test volume to local disk backup device after boot-time defragmentation & IFAAST engine run results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 4m 50s
Average Throughput 2,521.03 MB/min
Session Status Finished

22. BrightStor ARCserve Backup 11.5: Test volume to network disk backup device before defragmentation results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 16m 46s
Average Throughput 726.50 MB/min
Session Status Finished

23. BrightStor ARCserve Backup 11.5: Test volume to network disk backup device after defragmentation results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 16m 18s
Average Throughput 747.30 MB/min
Session Status Finished

Appendix continued:

24. BrightStor ARCserve Backup 11.5: Test volume to network disk backup device after boot-time defragmentation results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 9m 24s
Average Throughput 1,295.86 MB/min
Session Status Finished

25. BrightStor ARCserve Backup 11.5: Test volume to network disk backup device after IFAAST engine run results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 11m 24s
Average Throughput 1,068.85 MB/min
Session Status Finished

26. BrightStor ARCserve Backup 11.5: Test volume to network disk backup device after boot-time defragmentation & IFAAST engine run results (excerpted from the backup log):

Source. I:
Total Directories 18
Total File (s) 99994
Total Data Processed . . . 11,869.30 MB
Elapsed Time 5m 7s
Average Throughput 2,381.43 MB/min
Session Status Finished

27. Backup Exec 10d: Test Volume to network tape backup device after defragmentation results (excerpted from the backup log):

Set status : Completed
Resource name : [\\SYS332K3ENTSP1\I:](#)
Agent used : Yes
Start time : Tuesday, May 23, 2006 12:22:49 PM
End time : Tuesday, May 23, 2006 12:32:05 PM
Files : 99,998
Directories : 21
Skipped files : 0

Corrupt files : 0
Files in use : 0
Byte count : 12,461,503,885 bytes
Rate : 1,942.00 MB/Min

END

Diskeeper Corporation
7590 N. Glenoaks Blvd.
Burbank, CA 91504
800-829-6468
www.diskeeper.com



© 2006 Diskeeper Corporation. All Rights Reserved. Diskeeper, the Diskeeper Corporation logo and I-FAAST are trademarks or registered trademarks of the Diskeeper Corporation in the United States and/or other countries. Backup Exec is a trademark owned by the Symantec Corporation. BrightStor ARCserve Backup is a trademark owned by Computer Associates. Windows is a registered trademark of the Microsoft Corporation. Pentium is a trademark owned by the Intel Corporation. All other trademarks are the property of their respective owners.